



US006078953A

**United States Patent** [19][11] **Patent Number:** **6,078,953****Vaid et al.**[45] **Date of Patent:** **Jun. 20, 2000****[54] SYSTEM AND METHOD FOR MONITORING QUALITY OF SERVICE OVER NETWORK**

[75] Inventors: **Aseem Vaid**, San Jose; **Atul Pandit**, Campbell; **Gregory Rakoshitz**, Palo Alto; **Sanjeev R. Putta**, Campbell, all of Calif.

[73] Assignee: **Ukiah Software, Inc.**, Campbell, Calif.

[21] Appl. No.: **08/999,044**

[22] Filed: **Dec. 29, 1997**

**Related U.S. Application Data**

[60] Provisional application No. 60/067,857, Dec. 5, 1997, and provisional application No. 60/047,752, May 27, 1997.

[51] Int. Cl.<sup>7</sup> ..... **G06F 13/00**

[52] U.S. Cl. .... **709/223**; 709/232; 709/238; 370/235; 370/412

[58] Field of Search ..... 709/102, 103, 709/230, 238, 239, 240, 249, 223, 224, 232; 370/229, 235, 396, 401, 428, 230, 412, 413, 468, 469; 713/201

**[56] References Cited****U.S. PATENT DOCUMENTS**

5,251,209 10/1993 Jurkevich et al. .... 370/468

5,506,966	4/1996	Ban .....	709/234
5,606,668	2/1997	Shwed .....	713/201
5,675,573	10/1997	Karol et al. ....	370/230
5,696,764	12/1997	Soumiya et al. ....	370/395
5,742,596	4/1998	Baratz et al. ....	370/356
5,742,606	4/1998	Iliadis et al. ....	370/413
5,787,086	7/1998	McClure et al. ....	370/413
5,864,542	1/1999	Gupta et al. ....	370/257
5,913,041	6/1999	Ramanathan et al. ....	709/233

Primary Examiner—Viet D. Vu

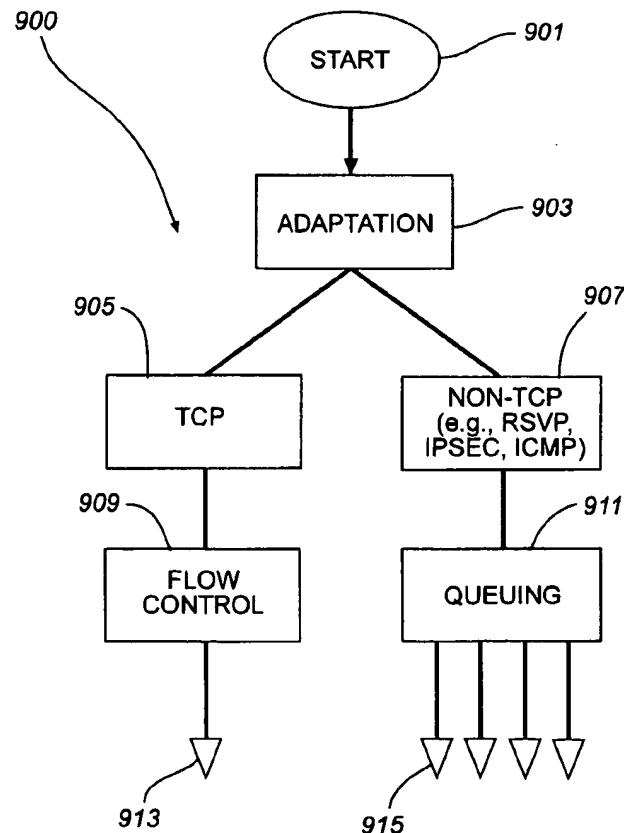
Attorney, Agent, or Firm—Townsend and Townsend and Crew LLP

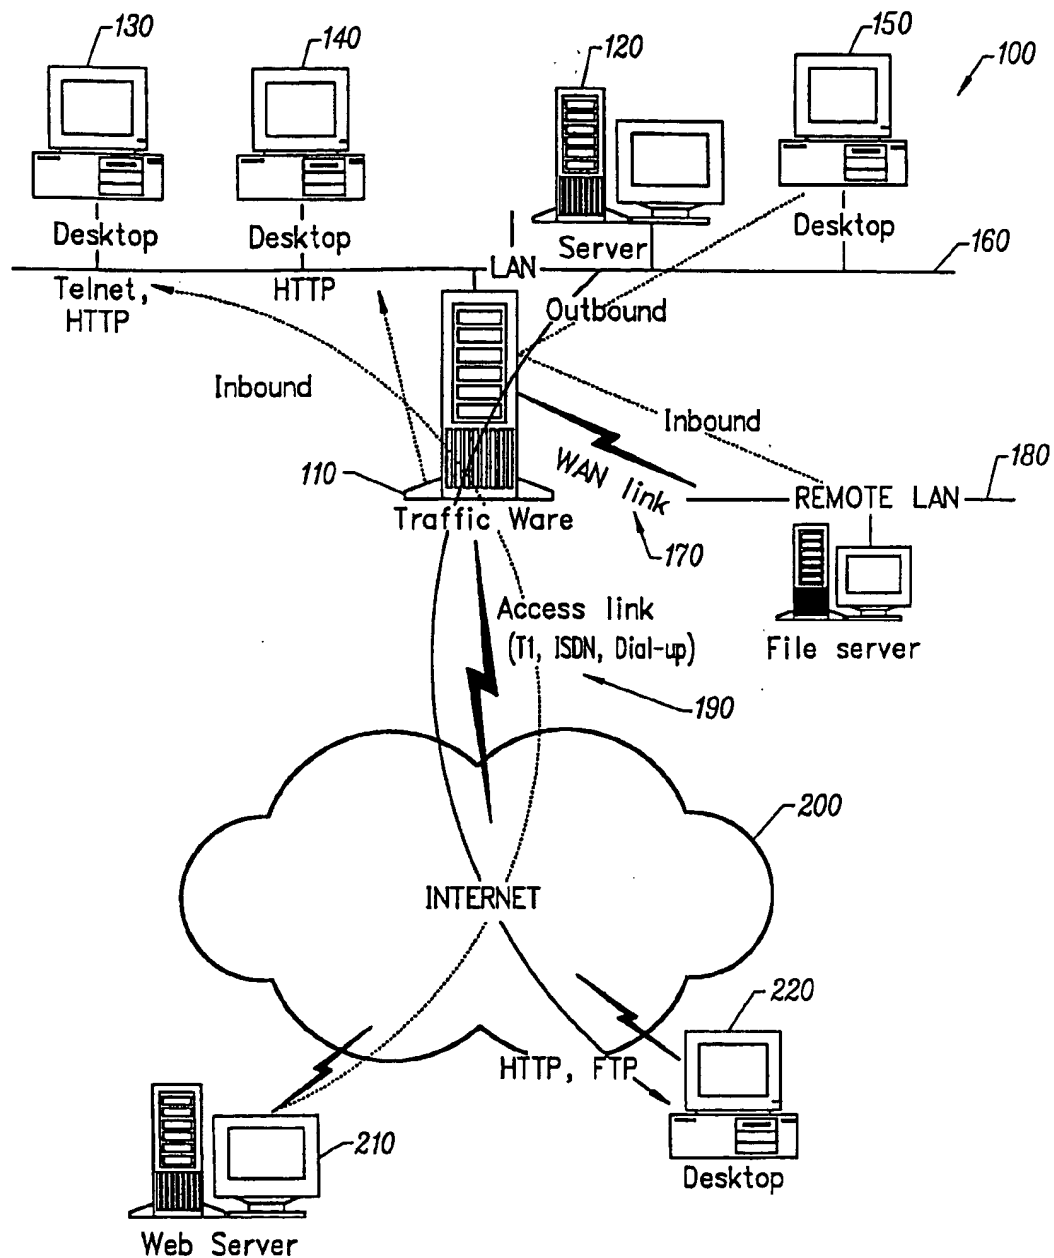
**[57] ABSTRACT**

A method and system (100) for controlling quality of service within a network of computers. The method includes a step of providing a network of computers, each being coupled to each other to form a local area network. The network of computers has a firewall server (110) coupled to the network of computers and a traffic management tool coupled to the firewall server. The method also includes implementing traffic control for incoming and outgoing information using a combination of flow control and queuing techniques.

network  
traffic mgmt  
tool

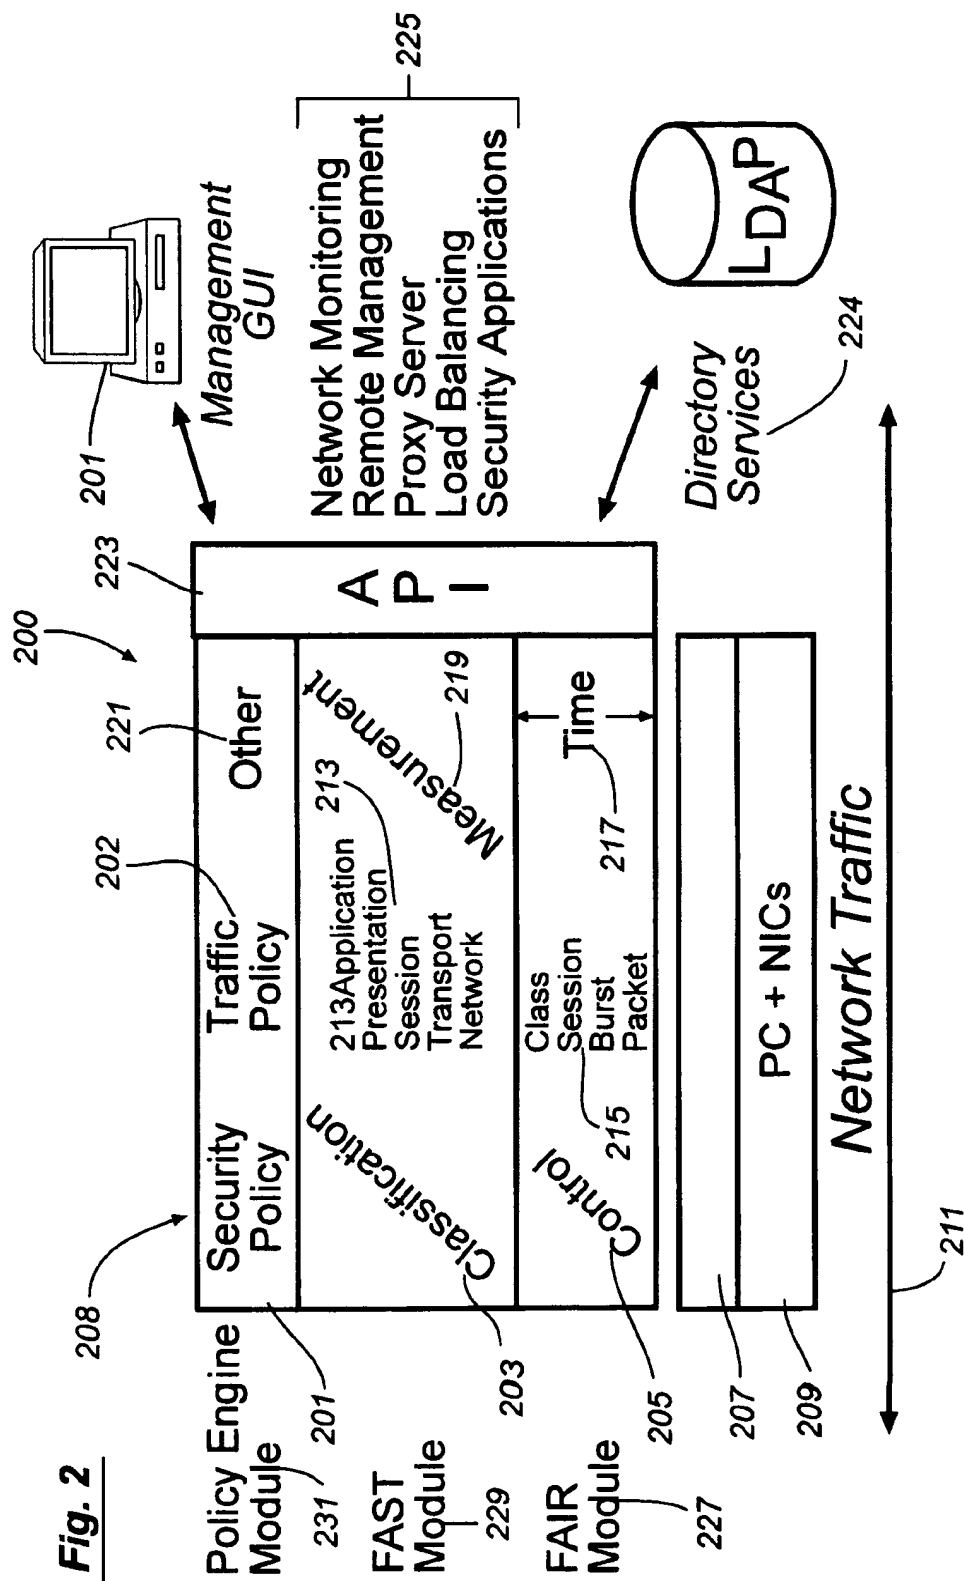
16 Claims, 9 Drawing Sheets

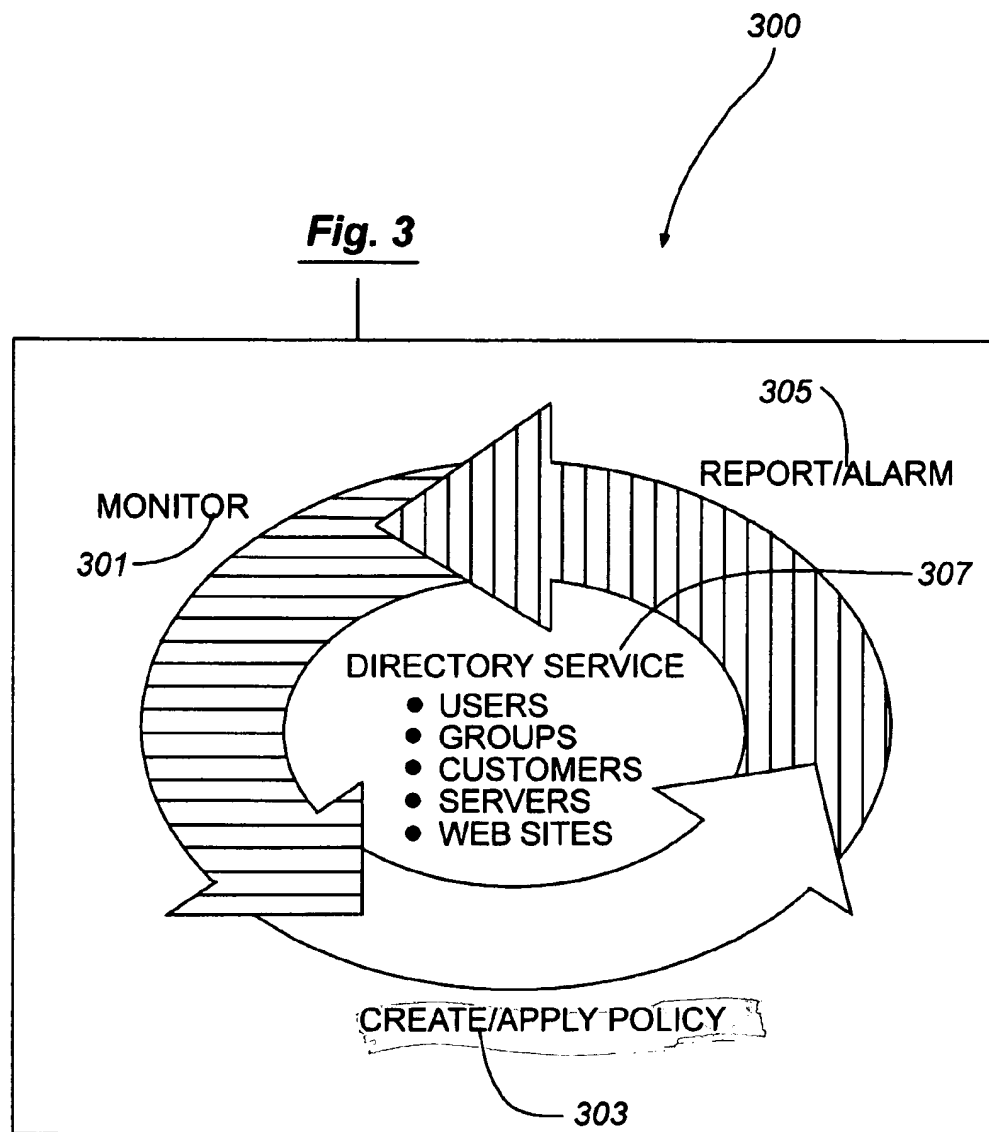


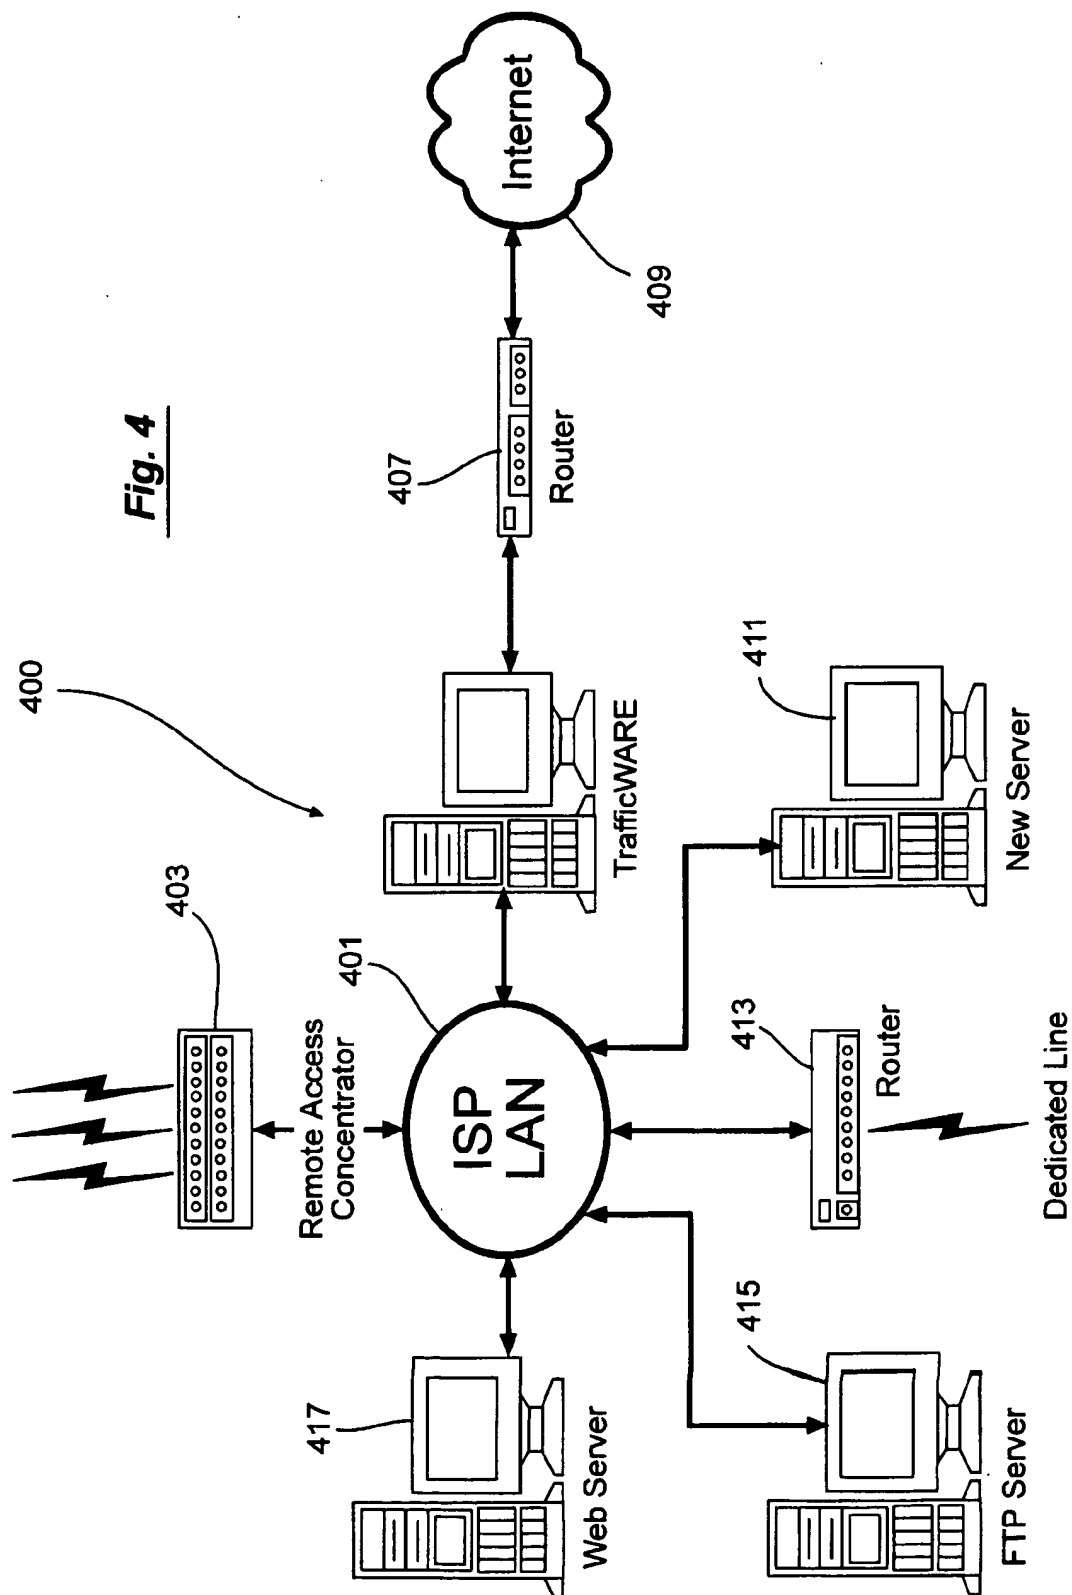


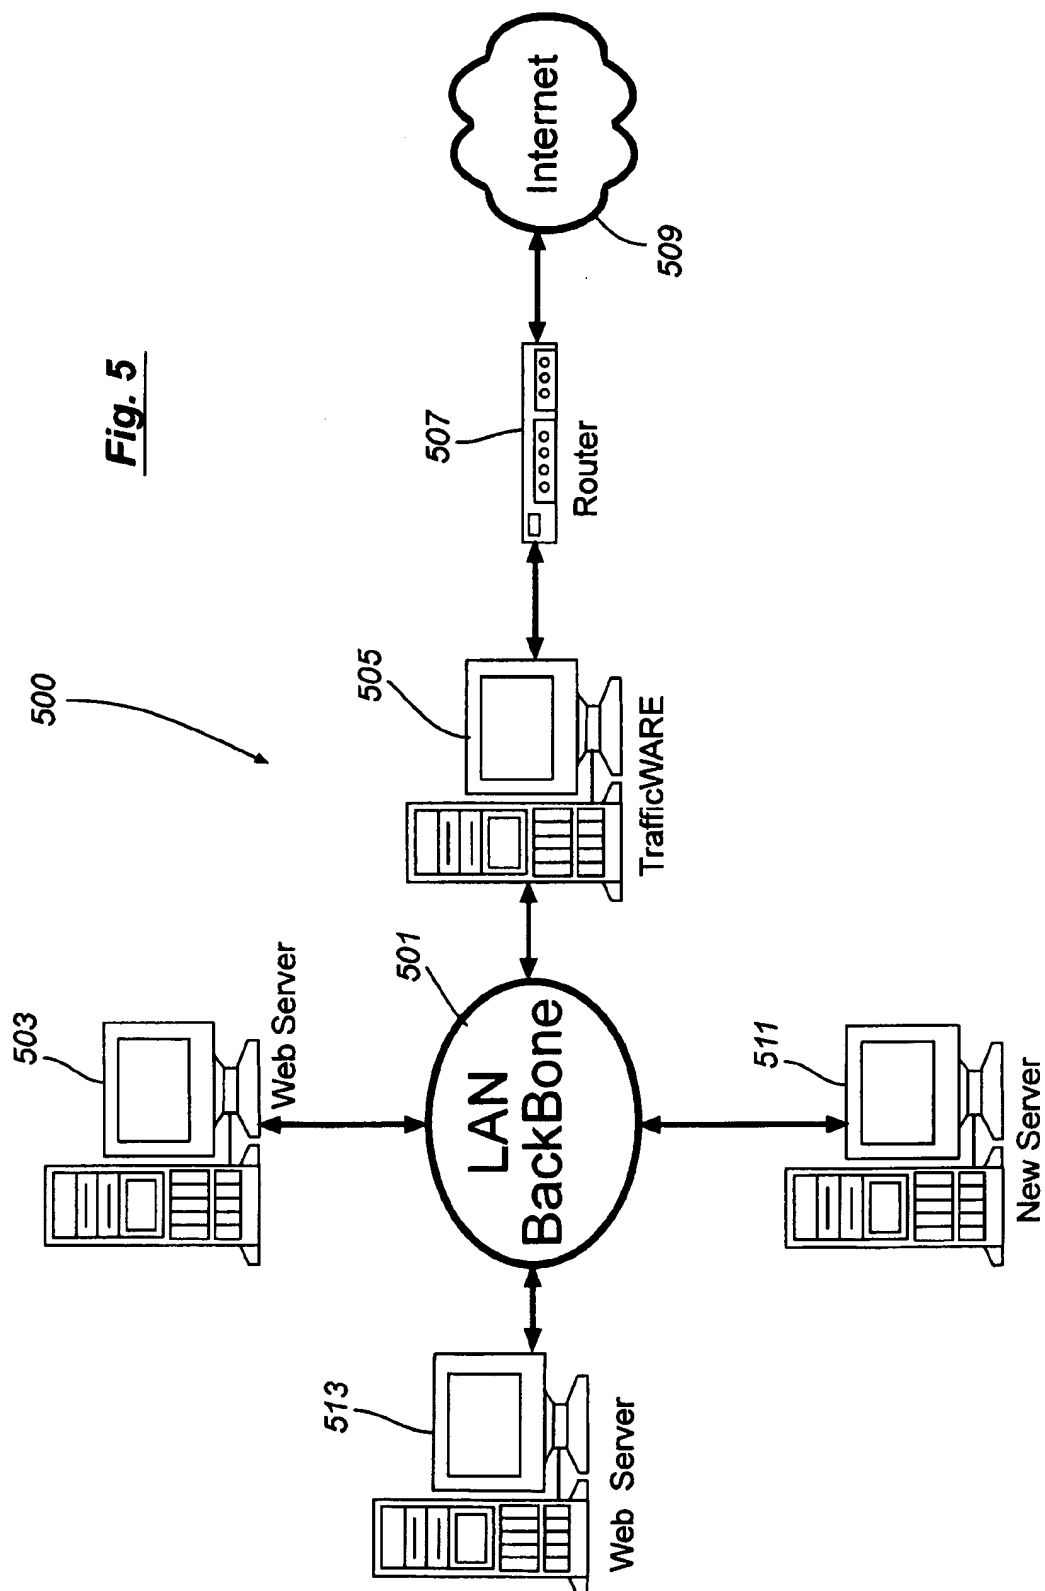
SCHEMATIC OF QOS SYSTEM DEPLOYED FOR INTERNET ACCESS  
AND A PRIVATE WAN

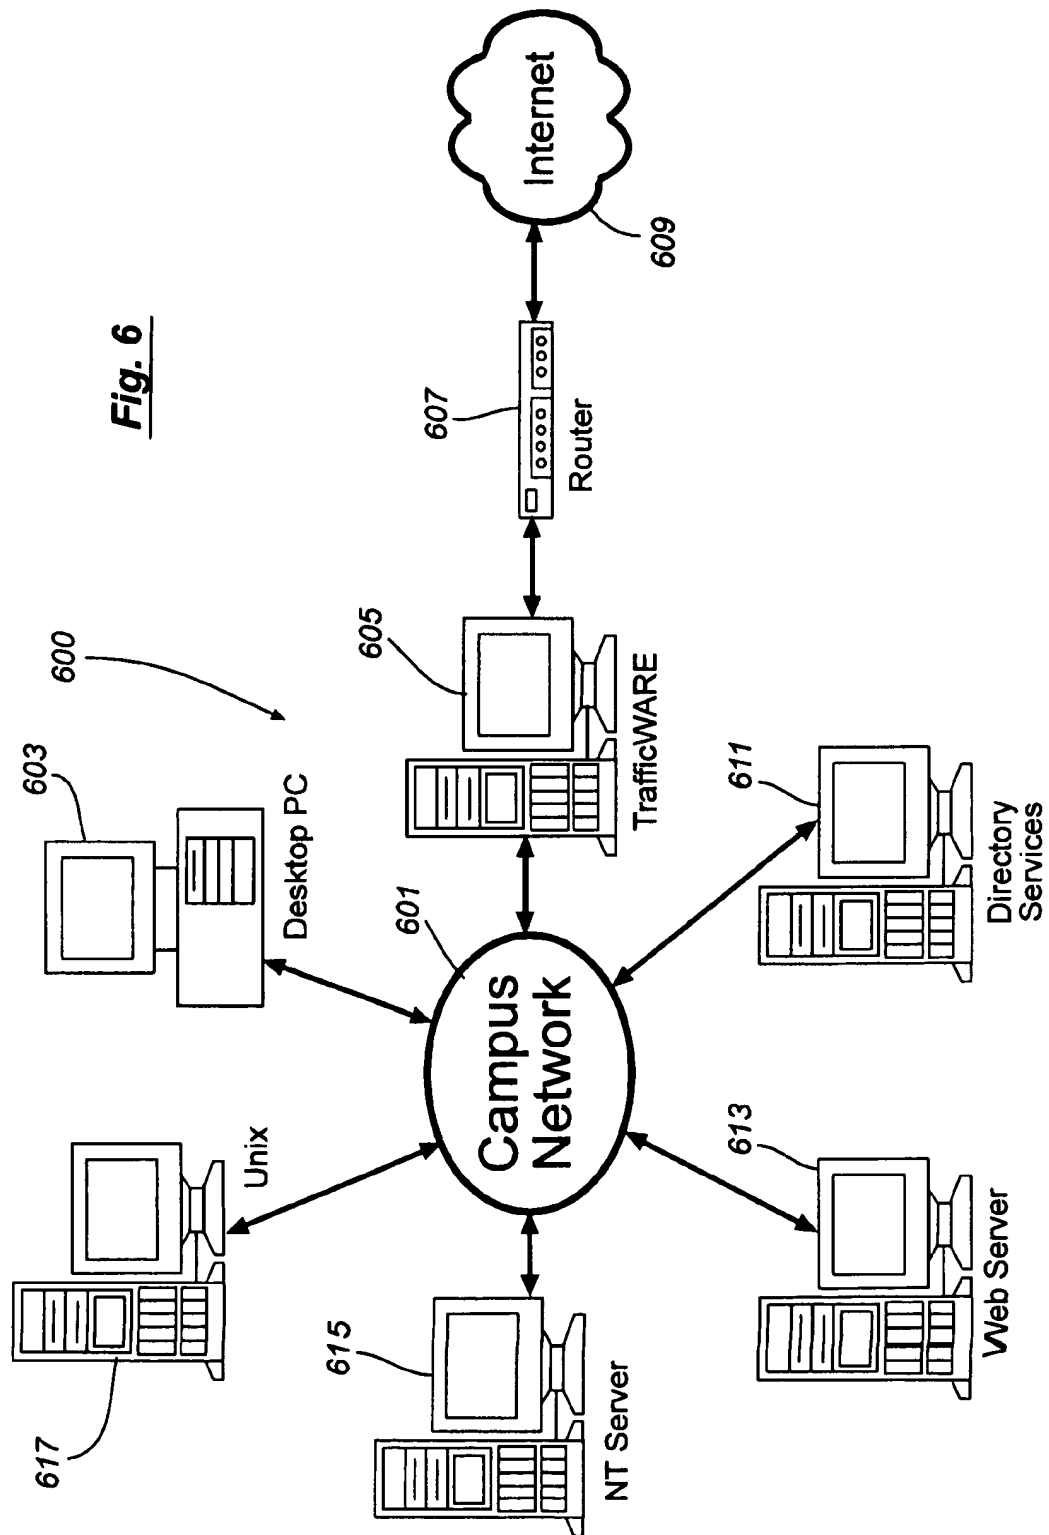
FIG. 1

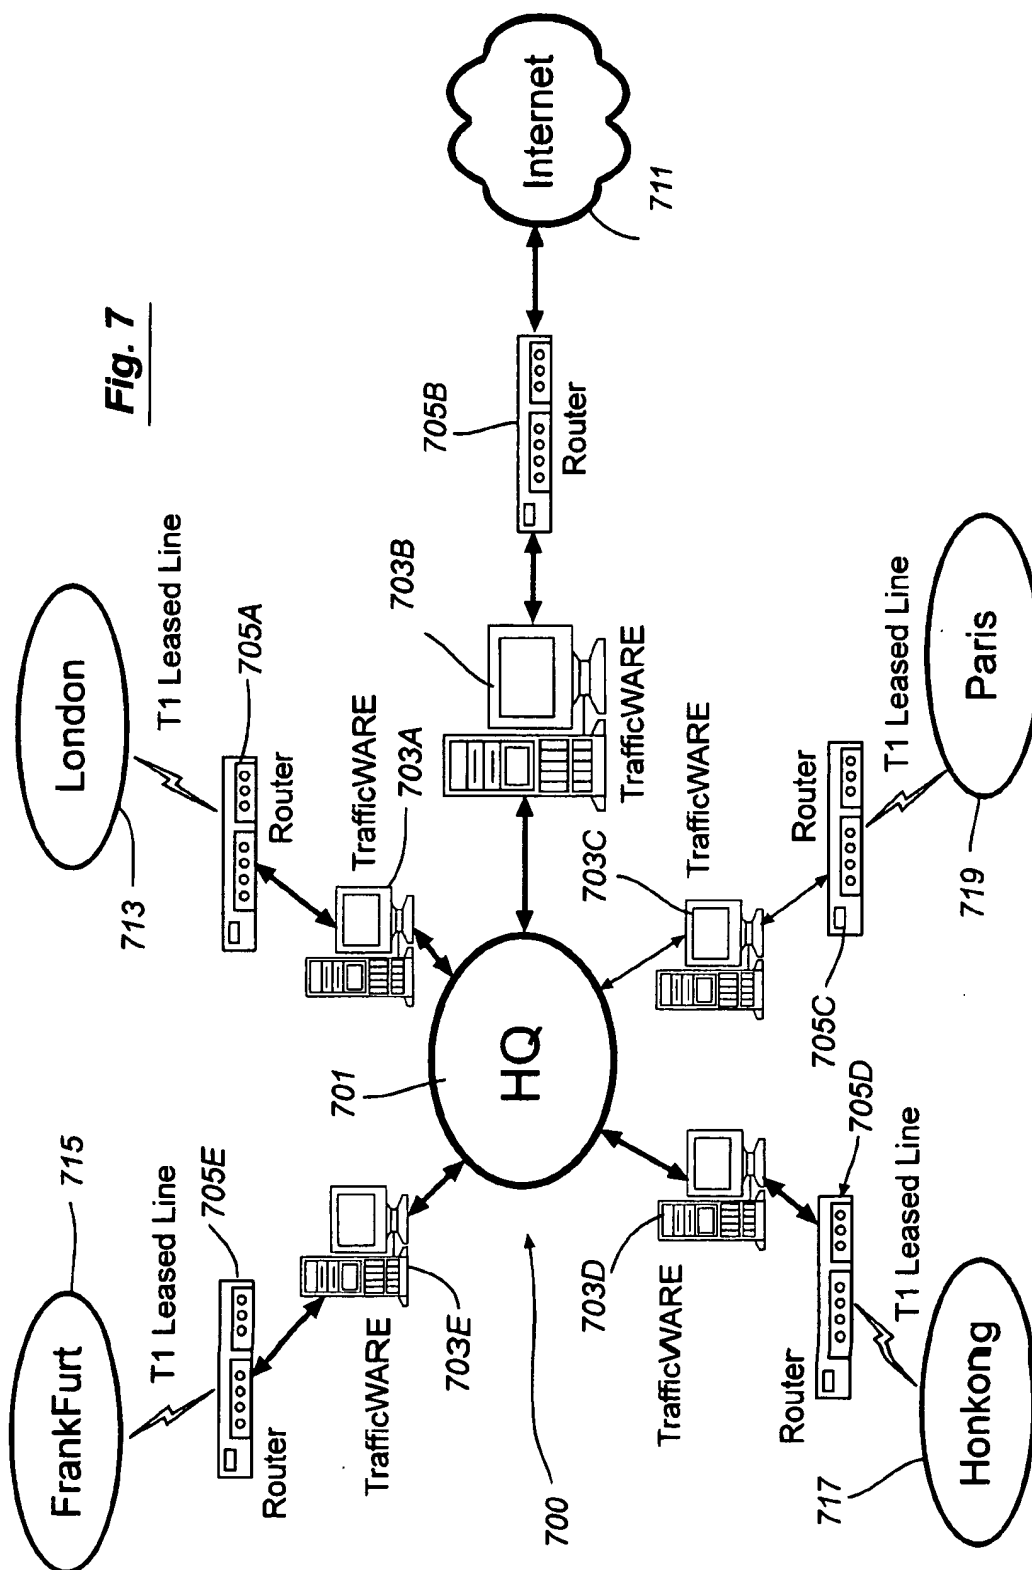


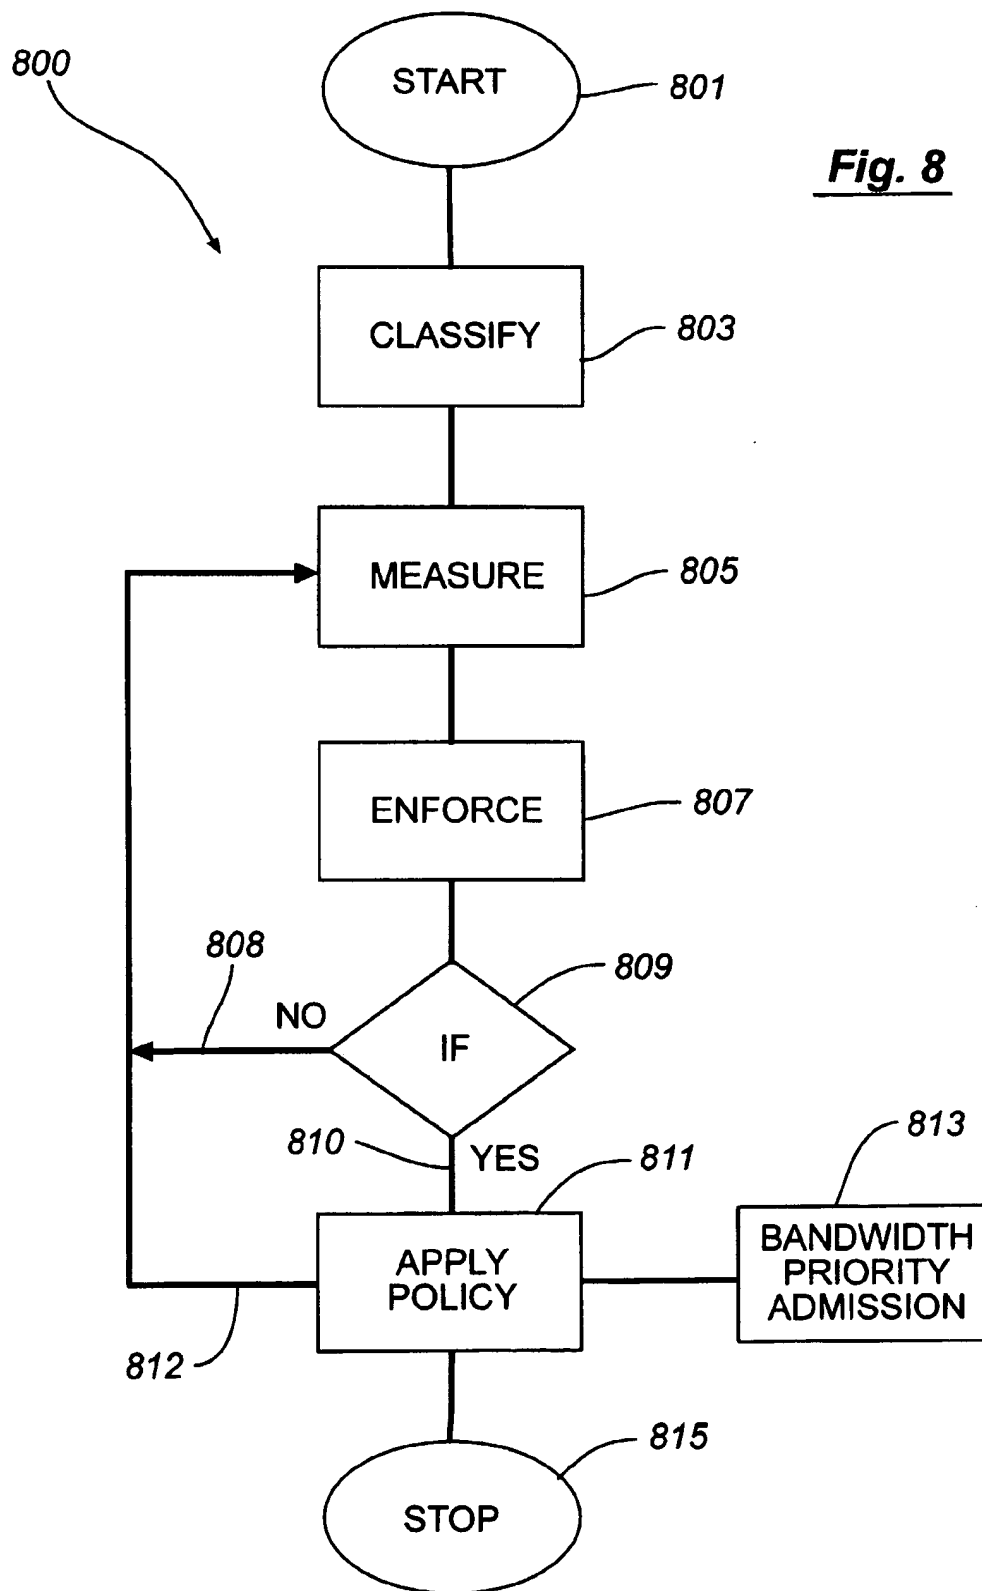


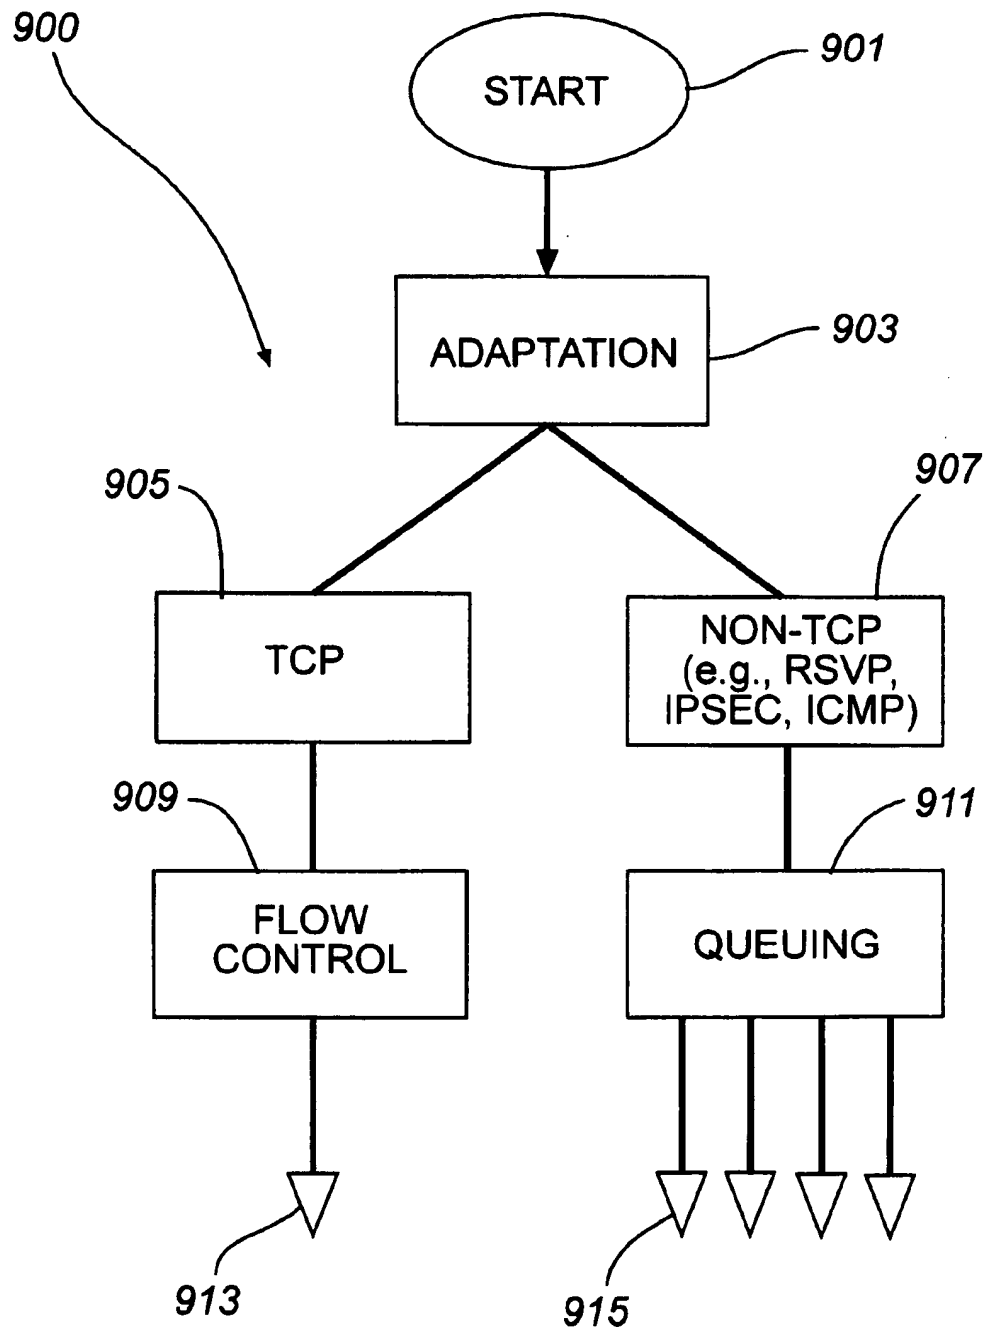












***Fig. 9***

## SYSTEM AND METHOD FOR MONITORING QUALITY OF SERVICE OVER NETWORK

### RELATED APPLICATIONS

This present application claims priority to U.S. Ser. No. 60/067,857 filed Dec. 5, 1997, and U.S. Ser. No. 60/047,752 filed May 27, 1997, which are both hereby incorporated by reference for all purposes.

### BACKGROUND OF THE INVENTION

The present invention relates to communication or telecommunication. More particularly, the present invention provides a technique, including a method and system, for monitoring and allocating bandwidth on a telecommunication network at, for example, a firewall access point. As merely an example, the present invention is implemented on a wide area network of computers or workstations such as the Internet. But it would be recognized that the present invention has a much broader range of applicability including local area networks, a combination of wide and local area networks, and the like.

Telecommunication techniques have been around for numerous years. In the early days, people such as the American Indians communicated to each other over long distances using "smoke signals." Smoke signals were generally used to transfer visual information from one geographical location to be observed at another geographical location. Since smoke signals could only be seen over a limited range of geographical distances, they were soon replaced by a communication technique known as telegraph. Telegraph generally transferred information from one geographical location to another geographical location using electrical signals in the form of "dots" and "dashes" over transmission lines. An example of commonly used electrical signals is Morse code. Telegraph has been, for the most part, replaced by telephone. The telephone was invented by Alexander Graham Bell in the 1800s to transmit and send voice information using electrical analog signals over a telephone line, or more commonly a single twisted pair copper line. Most industrialized countries today rely heavily upon telephone to facilitate communication between businesses and people, in general.

In the 1990s, another significant development in the telecommunication industry occurred. People began communicating to each other by way of computers, which are coupled to the telephone lines or telephone network. These computers or workstations coupled to each other can transmit many types of information from one geographical location to another geographical location. This information can be in the form of voice, video, and data, which have been commonly termed as "multimedia." Information transmitted over the Internet or Internet "traffic" has increased dramatically in recent years. In fact, the increased traffic has caused congestion, which leads to problems in responsiveness and throughput. This congestion is similar to the congestion of automobiles on a freeway, such as those in Silicon Valley from the recent "boom" in high technology companies, including companies specializing in telecommunication. As a result, individual users, businesses, and others have been spending more time waiting for information, and less time on productive activities. For example, a typical user of the Internet may spend a great deal of time attempting to view selected sites, which are commonly referred to as "Websites," on the Internet. Additionally, information being sent from one site to another through electronic mail, which is termed "e-mail," may not reach its destination in a timely

or adequate manner. In effect, quality of service or Quality of Service ("QoS") of the Internet has decreased to the point where some messages are being read at some time significantly beyond the time the messages were sent.

Quality of Service is often measured by responsiveness, including the amount of time spent waiting for images, texts, and other data to be transferred, and by throughput of data across the Internet, and the like. Other aspects may be application specific, for example, jitter, quality of playback, quality of data transferred across the Internet, and the like. Three main sources of data latency include: the lack of bandwidth at the user (or receiving) end, the general congestion of Internet, and the lack of bandwidth at the source (or sending) end.

A solution to decreasing data latency includes increasing the bandwidth of the user. This is typically accomplished by upgrading the network link, for example by upgrading a modem or network connection. For example, the network link may be upgraded to X2 modems, 56K modems, ADSL or DMT modems, ISDN service and modems, cable TV service and modems, and the like. Drawbacks to these solutions include that they typically require additional network service; they also require additional hardware and/or software, and further they require both the sender and receiver to both agree on using the same hardware and/or software. Although one user may have a much faster line or faster modem, another user may still rely on the same 1,200 kbaud modem. So, the speed at which information moves from one location to another location is often determined by the slowest information which is being transferred over the network. Accordingly, users of faster technology are basically going nowhere, or "running" nowhere fast, as is commonly stated in the network industry.

From the above, it is seen that a technique for improving the use of a wide area network is highly desirable.

### SUMMARY OF THE INVENTION

The present invention relates to a technique, including a method and system, for providing more quality to telecommunication services. More particularly, the present invention relates to quality of service management using a combination of flow control and queuing techniques to be applied to a flow of information at, for example, a firewall server. Quality of service is provided using a predominately software based bandwidth management tool in a specific embodiment.

In a specific embodiment, the present invention provides a system for providing quality of service within a network of computers. The present system has a network of computers, each being coupled to each other to form a local area network. A link is coupled to the network of computers. The system also has a traffic management tool coupled to the link, where the traffic management tool includes a flow control module and a queuing control module. The combination of these modules manage incoming and outgoing information to and from the link.

In an alternative embodiment, the present invention provides a novel method of managing information at a gateway site for improving quality of service to a network of computers. The method has a step of providing a flow of information to a bandwidth management tool disposed at a server. The flow of information is classified into at least a first portion and a second portion. The method applies flow control to the first portion of the information to be transferred to a first user location or a network link and applies queuing techniques to the second portion of the information

to be transferred to a second user location or a network link. The combination flow control and queuing of incoming and outgoing information to and from a communication link provides managed packet flow.

In a further aspect, the present invention provides a system for managing information at a gateway site for improving quality of service to a network of computers. The present system includes a variety of computer codes to form computer software or a computer program in computer memory. The program includes a first code directed to classifying a flow of information into at least a first portion and a second portion, and a second code directed to applying flow control to the first portion of the information to be transferred to a first user location or link. A third code is directed to queuing the second portion of the information to be transferred to a second user location or link. The system also includes additional codes for carrying out the techniques described herein and others.

Numerous advantages are achieved by way of the present invention over pre-existing or conventional techniques. In a specific embodiment, the present invention provides a single point or a single region to manage telecommunication traffic including directory services and bandwidth management. Additionally, in some, if not all embodiments, the present invention can be implemented at a single point of access such as a computer terminal or firewall, for example. Furthermore, the present invention can be predominately software based and can be implemented into a pre-existing system by way of a relatively simple installation process. Moreover, the present invention provides more valued applications and users with a more reliable and faster service. Less critical applications and users are provided with a service level that is appropriate for them in some embodiments. In most embodiments, available bandwidth in a system is fairly shared between equally prioritized users (e.g., no user can monopolize or "hog" the system). Still further, link efficiency improves due to overall congestion avoidance in most cases. Moreover, the present invention implements its traffic management technique using a simple and easy to use "rule" based technique. These and other advantages are described throughout the present specification, and more particularly below.

Further understanding of the nature and advantages of the invention may be realized by reference to the remaining portions of the specification, drawings, and attached documents.

#### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a simplified diagram of a system according to an embodiment of the present invention;

FIG. 2 is a simplified block diagram of system architecture according to an embodiment of the present invention;

FIG. 3 is a simplified diagram of a traffic management cycle according to an embodiment of the present invention;

FIGS. 4-7 are simplified diagrams of systems according to various embodiments of the present invention;

FIG. 8 is a simplified flow diagram of a rule-based control method according to the present invention; and

FIG. 9 is a simplified flow diagram of a control method using flow control and queuing control according to the present invention.

#### DESCRIPTION OF SPECIFIC EMBODIMENTS

An embodiment of the present provides integrated network service policies for firewall platforms. Specifically, the

present invention provides network or firewall administrators with the ability to implement policy-based schema for security and resource management on firewall platforms. In a specific embodiment, resource management includes Network Quality of Service (QoS) or "bandwidth" management techniques.

Network QoS occurs by managing the resources that serve network application traffic, for example. This typically includes the following resources: link bandwidth, application server bandwidth (CPU), and buffer space on generally all nodes (end-points, routers and gateways). Typically, data through-put is limited by the speed of Internet access links and by the server CPU capacity, and response time is determined by the number of hops in a route, physical length of the route, and extent of congestion in the route. There are various other factors that may affect QoS, such as the behavior of TCP/IP, severe congestion anywhere in the route, prioritization of traffic along the route, etc. To a network administrator, embodiments of the present invention provide discrimination of different traffic types and provide methods for enforcement of traffic flow by management to the above resources.

#### Definitions

In the present invention, it may assist the reader to understand some of the terms described herein. These terms have been briefly described below. These terms are merely examples and should not unduly limit the scope of the claims herein.

1. Traffic Management: A set of techniques or mechanisms including policies that can be applied in a network to manage limited network resources such as bandwidth and the like. These techniques are intended to improve overall network performance and efficiency. They are also intended to provide for more predictability and orderliness in the event of network congestion. The techniques should also isolate faults and provide visibility into performance problems. Additionally, they should meet the diverse user and application requirements as per an organization's business goals. Furthermore, traffic management is intended to increase the "goodput" traffic, based on the economic value and prevent the abuse of network resources.

2. Quality Of Service (QoS): The concept of Quality of Service (QoS) has been analyzed and discussed for a number of years in the networking industry, and was previously associated mostly with ATM technology. In a more generic sense, QoS describes the performance specifications that an application requires from the underlying infrastructure. Otherwise, the application will not run satisfactorily. Some applications are designed to run in a best-effort mode and can adapt to available bandwidth. Others are extremely sensitive to delays. Still others can produce large bursts in traffic which affects other applications while providing little perceptible improvements to the end-user. QoS specifications are closely associated with the expectations and perceptions of end-users and the organization they are part of.

3. Bandwidth: Bandwidth usually refers to maximum available bit rate for a specific application. In a specific embodiment, synchronous, interactive, and real-time applications, which are bandwidth-sensitive, can require minimum bandwidth guarantees, and can require sustained and burst-scale bit-rates. On the other hand, network administrators may want to limit bandwidth taken by non-productive traffic such as push technologies like PointCast and others. Even though bandwidth may be allocated for specified applications, it does not mean that these applications may be using that bandwidth. Therefore, a good policy should be to enforce when there is competition and demand.

*policies*

4. Latency: Latency generally refers to the delay experienced by a packet from the source to destination. Latency requirements are typically specified as mean-delay and worst case delay in some cases. Real-time audio/video applications such as, for example, DNS, HTTP, and TELNET are delay sensitive. Delay is a result of propagation delay, due to physical medium and queuing at intermediate nodes such as routers, gateways, or even servers. A certain portion of the delay can be controlled by how the queues are serviced at the intermediate nodes, and by controlling congestion at bottleneck points. Some examples of delay measures are packet round-trip delay and connection response time.

5. Jitter: Jitter generally refers to variation in delay (e.g., that is, the delay is not constant for all packets of a given flow) for a particular application. Real-time applications require a worst case jitter. Applications such as real-audio and video do some advanced buffering to overcome any variation in packet delays—the amount of buffering is determined by the expected jitter.

6. Packet Loss: Packet loss is a loss in a packet or a portion of packets that is generally caused by failure of network elements (e.g., routers, servers) to forward or deliver packets. Packet loss is usually an indication of severe congestion, overload of an element, or element failure (e.g., if a server is down). Even if the packet was not dropped but just delayed, protocols and applications can assume it was lost. Packet loss can cause application timeouts, loss of quality or retransmitted packets. Packet loss is usually specified as a rate (e.g., a real-time video application cannot tolerate loss of more than one packet for every 10 packets sent). Indirect results of packet loss may also be measured (e.g., connection retries or data retransmits).

7. Guarantees: An extreme example of a guarantee is to partition bandwidth so that it is not available to other entities. Guarantee also means a share of the resource, e.g., minimum bandwidth or maximum latency.

8. Best-effort: Best-efforts describes a service on best-effort basis but makes no guarantees.

9. Limits: Specific physical or theoretical limitation on a resource such as bandwidth. Resource utilization or admission is limited under certain conditions.

10. Priority: Level of importance for a specific user, application, or data. Create a priority scheme among different entities so that contention is resolved or service is provided.

The above definitions are merely intended to assist the reader in understanding some of the terms described herein. They are not intended, in any manner, to limit the scope of the claims. One of ordinary skill in the art would recognize other variations, modifications, and alternatives.

## System Overview

FIG. 1 illustrates a simplified system 100 according to an embodiment of the present invention. The system 100 is merely an illustration and should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. The present invention can be embodied as a TrafficWare™ firewall server 110 from Ukiah Software, Inc, but can be others. System 100 typically includes a file server 120, and a plurality of computers 130–150, coupled to a local area network (LAN) 160, and other elements. Firewall server 110 includes a typical connection to a wide area network (WAN) 170 and to a remote LAN 180 (such as an Intranet) and a typical network connection 190 to the Internet 200. Attached to Internet 200 are Web servers 210 and other computers 220.

As illustrated, computers such as computer 130, 140, and 210 communicate using any one or multiple application layer protocols such as Telnet, file transfer protocol (FTP), Hypertext transmission protocol (HTTP), and the like. Further, communication across WAN 170 and across network connection 190 implements transport layer protocols such as transmission control protocol (TCP), universal data protocol (UDP), and the like. LAN 160 and LAN 180 are preferably based upon network protocols such as Internet protocol (IP), IPX from Novell, AppleTalk, and the like. As shown in FIG. 1, network connection 190 may be accomplished using T1, ISDN, Dial-up, and other hardware connections. Computers 120–150 and 210–220 may be any suitable make or model of computer that can be coupled to a network. The system can also include a variety of other elements such as bridges, routers, and the like.

In an alternative specific embodiment, the present invention may be applied to a system with various links accessed in servicing a browser request at a remote web server. In this embodiment, a client could be dialing in via a 28.8 kbit dial up modem to a local Internet service provider (ISP), where the ISP may be connected to the Internet by a T1 link. A web server may be on a 10 Mbs Ethernet LAN, which is connected to another ISP via a 56 K frame relay. The web server's ISP may be connected to its carrier via a T3 line. The client ISP carrier and the server ISP carrier may both be connected by an ATM backbone or the like. Because of this asymmetry in this embodiment, any traffic management solution should take into account these variations including traffic speed and data format described above. Moreover, simply upgrading the capacity of a link, in the access path, may not offer a viable solution. This present embodiment may have the bandwidth requirements shown by way of Table 1, for example.

TABLE 1

Bandwidth Requirements		
Users	Bandwidth	Service Offered
Internet developers, individuals, international locations where bandwidth is expensive	28.8 to 56 Kbps	Dial-up services, ISDN
Small to medium-sized organizations with moderate Internet usage	56 Kbps to 1.5 Mbps	Fractional T1, frame relay
Medium sized organizations with many moderate users, smaller organizations requiring huge amounts of bandwidth	1.5 Mbps	Dedicated T1 circuit
Standard bandwidth for Ethernet-based LANS	10 Mbps	Ethernet, token ring (4 Mbps or 16 Mbps)
Bandwidth usage for large organizations or	45 Mbps	Dedicated T3 circuit

TABLE 1-continued

Users	<u>Bandwidth Requirements</u>	
	Bandwidth	Service Offered
Internet backbones Huge bandwidth LAN backbone usage for medium to large organizations (hundreds or thousands of users)	100 to 1,000 Mbps	Fast Ethernet, gigabit Ethernet

As shown above, there exist a large number of diverse applications and protocols that are widely used and have their own performance requirements. For example, applications such as mail (e.g., SMTP) and news (e.g., NNTP) are not interactive and are therefore not sensitive to delay. On the other hand, applications such as real-time conferencing are extremely sensitive to delay but not to packet loss. Applications such as TELNET and DNS do not utilize significant bandwidth, but are sensitive to delay and loss. Conversely, applications such as FTP consume a great deal of bandwidth but are not that sensitive to delay. Generally, network applications can be categorized as:

1. Interactive (e.g., delay sensitive) versus non-interactive (e.g., delay tolerant);
2. Bandwidth intensive (bulk data) versus non-bandwidth intensive; and
3. Bursty versus non-bursty.

These categories are merely illustrative and should not limit the scope of the claims herein. Additionally, some application requirements are dependent on the context of use and the nature of data being accessed. Such applications can be described as being nominally interactive or nominally bandwidth intense. This means the description applies to many but not all the situations in which they are used.

As merely an example, Table 2 provides some illustrations for these categories.

TABLE 2

<u>Application Spectrum</u>	
Application Class	Examples
Low-bandwidth, delay sensitive, highly interactive	DNS, PING, TELNET, CHAT, COLLABORATION
High bandwidth, delay sensitive	Real-time audio and video
High Bandwidth, nominally interactive	Web service requests, file downloads
Non-interactive	Mail and news

As shown in Table 2, low-bandwidth, delay sensitive, and highly interactive applications include, among others, DNS, PING, TELNET, CHAT, COLLABORATION. High bandwidth and delay sensitive applications including at least real-time audio and video. Additional applications for high bandwidth and nominally interactive, or non-interactive have also been shown. Again, these applications are merely provided for illustration and should not limit the scope of the claims herein.

The present invention can also be used with a number of various files. For example, a number of common applications, such as FTP and HTTP, can handle a wide variety of files. The file types being transferred and downloaded place different demands on the underlying infrastructure. Index and HTML files take up limited bandwidth but have very mundane contents. On the other hand, GIF, JPEG, and MPEG, RA and AVI files take up a lot more bandwidth but provide a rich multimedia experience to the end-user. In fact, push technologies such as PointCast basically download rich-multimedia bandwidth-intensive files.

The present invention can also be used with a variety of user requirements. For example, networks are facing an explosion in the number of (inter) networked applications and data accessible through them. Network resources are increasingly being used for a wide variety of purposes, ranging from business critical to personal. This means that policies must ensure that scarce resources (e.g., Internet bandwidth) are utilized with the goal of maximizing the returns to the organization. These benefits can come from direct revenue generating activities or from improved productivity (or reduced loss of productivity). As shown in Table 3, for example, at a mythical company called "Shebang Software Inc." the highest bandwidth priority has been allocated to technical support. However, there is no hard and fast rule. As with security policies, decisions should be consistent with the needs of the organization.

TABLE 3

<u>Shebang Software User Priorities</u>		
Users	Application Class	Reasons
Technical support	Mission critical	Needs most bandwidth to deal with customers who need assistance
Sales and marketing	Critical	Needs bandwidth to deal with potential customers. Answer inquiries, make quotes, transmit multimedia presentations
Upper management and middle management, administrative	Casual	Needs bandwidth to perform tasks necessary to run the business
Development and manufacturing	Personal	Needs bandwidth to send e-mail, subscribe to Push technologies

The present invention takes into account, in one or more embodiments, the factors which are described specifically above. Although the above has been generally described in terms of a specific type of information, other types of information on a network can also be used with the present invention. Additionally, the present invention has been described in general to a specific system. For instance, the present bandwidth management tool can be applied at a network's Internet access link. Alternatively, the present tool can be applied to a private WAN link to a remote corporate site or an access to a server farm (e.g., a group of servers located in a special part of the network close to an access link, e.g., in a web hosting environment). Alternatively, the present invention can be applied to key servers (e.g., database/web server) within an organization servicing internal and/or external users. Furthermore, the present bandwidth management tool can be applied to any combination of the above or the like.

FIG. 2 is a simplified block diagram 200 of details of system architecture according to an embodiment of the present invention. The block diagram is merely an illustration and should not limit the scope of the claims herein. The architecture includes a variety of layers that each interface to each other as depicted by the layers. The system includes a network layer 211, which interfaces to incoming and outgoing information to the network. The network can be one of a variety including, among others, Ethernet and Token Ring. A physical layer 209 is disposed above the network layer 211. The physical layer can be personal computers, which are commonly called PCs, or network interface computers, which are commonly called NCs, or alternatively workstations. As merely an example, a personal computer can be an IBM PC compatible computer having a 586-class based microprocessor, such a Pentium™ from Intel Corporation, but is not limited to such a computer or processor. An operating system ("OS") is used on the computer such as WindowsNT™ from Microsoft Corporation, but can also be other OSs. The system is also coupled to a graphical user interface ("GUI") 201 and is coupled to directory services such as, for example, LDAP, but can be others. A detailed discussion of directory services is described in U.S. application Ser. Nos. 08/998,100, 08/998,313, and 08/999,096 which are commonly assigned, and hereby incorporated by reference for all purposes.

Directory services 224 and GUI 201 couple to an application programming interface ("API") 223. The API is coupled to a traffic management or bandwidth management tool 208 with at least three modules, including a policy engine module 231, a FAST module 229, and a FAIR module 227, which will be discussed in more detail below, but is not limited to these modules. The bandwidth management tool 208 can be predominantly software based and is substantially free from any significant hardware or software changes in the network. In a preferred embodiment, the bandwidth management tool 208 can be loaded onto a server without any changes to hardware. In an alternative preferred embodiment, the tool can install, configure, and operate on a conventional IBM compatible PC running and operating system such as, for example, Windows NT, but can be others. The tool can be deployed at any appropriate point in the network data path. The tool can also be stand-alone at the WAN access point (e.g., behind the Internet access router or behind a firewall), with a conventional firewall or with an NT based proxy/caching server or application server (e.g., a Web server).

Tool 208 performs incoming and/or outgoing management of information over the network of computers. In a

specific embodiment, traffic management tool 208 performs inbound and outbound monitoring and control of flows by application, source address, destination address, URL, time of day, day of week, day of month, and other variations. In a specific embodiment, tool 208 also monitors, controls, and produces reports and alarms, which can enhance a whole spectrum of traffic monitoring and control activities ranging from bandwidth/latency control to capacity planning.

In a specific embodiment, the bandwidth management tool adapts to "real" changes on any pre-existing networking system. For example, network infrastructure management involves a continuous process of monitoring, reporting, and deploying changes to match network growth or changing needs in a growing office, for example. These changes exist at various levels and time scales. As merely examples, the network changes can be to enforce a QoS Policy for a critical service, add WAN bandwidth, segment the network, upgrade a router, choose a guaranteed service level for a web site (e.g., user's own web site), or notify "Mr. Hog" (i.e., a user occupying too much bandwidth) that he should schedule his large personal downloads at more prudent times such as late at night, for example.

#### Bandwidth Management Process

The bandwidth management tool can employ these changes using, for example, the process shown in FIG. 3. This process is merely an illustration and should not limit the scope of the claims herein. As shown, FIG. 3 is a simplified diagram 300 of a traffic management cycle according to an embodiment of the present invention. The traffic management cycle is depicted as a continuous cycle, which includes a monitoring phase 301, a creating/applying policy phase 303, and a reporting/alarms phase 305, but is not limited to these cycles. That is, these cycles can be separated or combined depending upon the application. By way of this cycle, the tool can adapt to any changes to the networking system according to the present invention.

In an aspect of the present invention, the present tool can monitor and control activities at various times, e.g., seconds, days, weeks, months, years. Some details with regard to these control activities are shown below under the headings.

#### 1. Second to second

The tool provides second to second time scale monitoring and control of incoming and outgoing traffic over the network. As merely an example, the tool ensures that critical or more important traffic gets a right of way during traffic bursts and provides bandwidth enforcement. Multiple users of the network at a specific time can cause the traffic burst. Alternatively, multiple sessions on the network at a specific time can cause the traffic burst. Once the traffic burst is detected, the tool has a control device, which provides bandwidth enforcement to ensure that the more important traffic gets through the network.

#### 2. Day to day

The tool provides day to day time scale monitoring and control of incoming and outgoing traffic over the network. As merely an example, the tool manages time of day congestion, and responds to intermittent problems or perceived problems. The tool generally deals with problems or limitations that are very specific and isolated to particular users or particular services at particular times that need to be tracked down quickly.

#### 3. Week to week

The tool provides week to week time scale monitoring and control of incoming and outgoing traffic over the network. The tool analyzes traffic usage performance patterns, what services or hosts are active on the network, and troubleshoots chronic problems. In particular, the tool looks

at aggregates, such as a particular segment of the network, and compares Websites or compares groups of users for usage of bandwidth and frequency of usage.

#### 4. Longer term activities

The tool provides long term time scale monitoring and control of incoming and outgoing traffic over the network. The tool implements changes in organizational priorities, in billing. The tool also provides service for new applications as they are introduced, and provides for capacity planning for network resources. The present tool can also be used with network stress testing tools to obtain detailed analysis of flows and traffic behavior with/without policy enforcement before a new application is deployed to change the network infrastructure.

Based upon the above description, the present tool can be used to monitor and control incoming and outgoing traffic over a variety of time frequencies. The time frequencies include second by second, day to day, or long term, and combinations thereof, depending upon the application. Of course, the time frequency used depends upon the particular network and applications.

FIGS. 4-7 are simplified diagrams of systems according to various embodiments of the present invention. These diagrams are merely illustrations and should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, alternatives, and modifications. These systems show various deployment scenarios according to the present invention.

#### 1. Internet Service Provider (ISP)

FIG. 4 is a simplified diagram 400 of the present tool in an ISP environment according to the present invention. The diagram 400 includes a variety of elements such as an ISP LAN 401, which is coupled to network elements including a remote access concentrator 403, a web server 417, an FTP server 415, a router 413, a news server 411, and others. The tool 405 is coupled between the ISP LAN and router 407, which is connected to the Internet 409. In this embodiment, the ISP is providing a number of services to its customers and the present tool sits by the Internet link and manages inbound and outbound traffic.

#### 2. Web Hosting Deployment

FIG. 5 is a simplified diagram 500 of the present tool in a web hosting environment according to the present invention. The diagram 500 includes a variety of elements such as a LAN Backbone 501, which is coupled to network elements including web servers 503, 511, 513, and others. The present tool 505 is coupled between LAN 501 and router 507, which is connected to the Internet 509. In the present embodiment, the tool is being used to manage inbound and outbound traffic between some Websites and the Internet. In a specific embodiment, most of the data being transmitted is multimedia-based, but is not limited as such data.

#### 3. End-User Deployment

FIG. 6 is a simplified diagram 600 of the present tool in a campus environment according to the present invention. The diagram 600 includes a variety of features such as a campus network 601, which is coupled to network elements such as a desktop PC 603, a UNIX computer 617, an NT Server 615, a web server 613, directory services 611, and others. A bandwidth management tool 605 is coupled between campus network 601 and router 607, which is coupled to Internet 609. In this embodiment, a LAN or WAN supports a number of different setups and configurations, which are compete for bandwidth to access the Internet. The present tool acts as an arbitrator for implementing rules, enforcing policies, and setting admissions for classes, as well as perform other acts.

#### 4. Private WAN

FIG. 7 is a simplified diagram 700 of the present tool deployed for a large corporation that has an Intranet as well as an Internet. The diagram 700 includes a variety of elements or "children" such as a connection to Frankfurt 715, a connection to London 713, a connection to Hong Kong 717, and a connection to Paris 719. Each connection or child includes a router 705A, E, D, C, and the present tool 703A, E, D, C, which is coupled between the router and the hub ("HQ"). In a WAN-based environment, for example, HQ 701 is the hub that handles a number of independent systems (e.g., Frankfurt, London, Hong Kong, Paris), which can be LAN-based. In this embodiment, the present tool 703B also sits by the Internet 711 and is used to allocate bandwidth between the competing children, e.g., Frankfurt, London, Hong Kong, Paris. Router 705B is coupled between tool 703B and Internet 711.

Although the above descriptions have been made in terms of deploying the present tool in selected environments, the present tool can also be deployed in other environments. For example, the present tool can be deployed in any combination of the above. Alternatively, the present tool can be deployed in any portion of the above environments. Of course, the type of environment used by the present tool depends highly upon the application.

In a specific embodiment, the tool provides an easy to use interface or graphical user interface ("GUI") for performance monitoring and profiling (e.g., accounting). Profiling can be based on active services, clients and servers, among other parameters. Additionally, profiling of the network can be started as soon as the tool is installed into the server of the network. Accordingly, the tool provides immediate accounting and service measurement on a variety of QoS measures.

In a specific embodiment, the present tool generally uses two mechanisms to implement efficient traffic monitoring and traffic control. These mechanisms include processes performed by the FAST module and the FAIR module, which are shown in FIG. 2, for example. Additionally, the present tool uses a policy engine module 231, which oversees the FAST module 229 and the FAIR module 227. Some details of these modules are described as follows.

#### 1. FAST Module (Flow Analysis and Session Tagging)

The FAST module generally provides for monitoring of incoming and outgoing information to and from the network or link. Flow Analysis and Session Tagging ("FAST") implements rich, application level traffic classification, and measurement. This operation is accomplished without introducing slow data paths to minimize latency and maximize overall throughput of traffic through the tool management engine. As shown in the Fig., the FAST module provides for classification 203 of information such as parameters 213 including application, presentation, session, transport, and network. The FAST module also provides for measurement 219 of various parameters. The FAST module is coupled to the API.

#### 2. FAIR Module (Flow Analysis and Intelligent Regulation)

The FAIR module generally implements traffic control and manages bandwidth of incoming and outgoing information to and from the network or link. Flow Analysis and Intelligent Regulation ("FAIR") implements traffic control based on a combination of flow control and queuing algorithms. FAIR's objective provides inbound and outbound traffic management for meaningful time intervals, reducing the load on packet classifiers and packet schedulers. The FAIR module controls 205 incoming and outgoing informa-

tion to and from the network. Additionally, the FAIR module controls 205 by parameters 215 such as class, session, burst, packet, and others. The FAIR module also controls time 217 of allocating bandwidth for these parameters. The FAIR module is coupled to the API.

### 3. Policy Engine Module

The policy engine module 231 oversees the FAST and FAIR modules. The engine module also interfaces with the API. In an embodiment, the policy engine module includes a security policy 201, a traffic policy 202, and other policies 221. The security policy provides parameters for securing the present tool. The traffic policy defines specific limitations or parameters for the traffic.

Some definitions about the various modules have been described above. These definitions are not intended to be limiting. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Additionally, the modules described are generally provided in terms of computer software. Computer software can be used to program and implement these modules, as well as others. The modules can be combined or even separated, depending upon the applications. Functionality of the modules can also be combined with hardware or the like. In a specific embodiment, the present modules are implemented on an WindowsNT™ operating system, which has been developed by Microsoft Corporation. Of course, other operating systems can also be used. Accordingly, the present modules are not intended to be limiting in any manner.

In an embodiment, the present tool can be configured based upon at least the following components—traffic classes, traffic policies, traffic rules, and traffic entities. Some information about these components are described below.

#### 1. Traffic Classes

The present tool identifies data flows at a network site based on traffic classes. A traffic class is any combination of the following, but is not limited to these:

- IP address, subnet, network, netgroup, or range of source or destination;
- URL of the sender or group of URLs;
- Service (e.g., HTTP, FTP) or groups of services;
- FTP and HTTP, file types can be selected as well;
- Time of day, day of week/month; and
- Inbound and outbound information.

As shown above, traffic classes are directional. Traffic classes configured for inbound traffic are managed separately from traffic classes configured for outbound traffic. For example, the present tool may decide to guarantee a minimum bandwidth to critical traffic so that it is not affected by congestion from large downloads. Additionally, the present tool may want to monitor Push traffic for a while and then choose to limit it if it is perceived as a problem. Traffic classes can also be for measurement only or for control and measurement in some embodiments. These are merely examples and should not limit the scope of the claims herein.

#### 2. Traffic Policies

Traffic policies are generally mechanisms used to control the traffic behavior of specific classes. In an embodiment, the present tool can configure policy properties which provide, for example:

- Bandwidth guarantees—granting classes a minimum bandwidth in the presence of congestion or competition;
- Bandwidth limits—establishing a limit on the total bandwidth used by the class;
- Setting priorities—establishing a priority order for bandwidth limiting or servicing traffic from a class. (That is,

high priority classes are serviced first and are affected the least during contention for bandwidth. Lower priority classes are serviced in order of priority and may be more affected by congestion or contention);

Admission control—establishing conditions under which a new network session or service request is admitted or not admitted. (This kind of policy establishes a broad bandwidth control or service quality for sessions already admitted).

As shown, the present invention provides policies such as bandwidth guarantees, bandwidth limits, setting priorities, admission control, and others. It may assist the reader in understanding some of the terms used in the policies by drawing an analogy with a geographical highway for automobiles. For example, bandwidth relates to how fast one can go (e.g., fast or slow lane) once a user has entered the stream of traffic on the highway. That is, the physical limit for speed in the specific lane chosen. Priority is analogous to how quickly the user is able to enter the highway and move into a designated lane, and how often the user may have to temporarily give way to other vehicles during the drive. Admission control is analogous to the metered lights at the entrance of the freeway where one is made to wait under certain conditions. Of course, depending upon the applications other analogies can be used to explain the policies. Additionally, the policies are merely examples and should not limit the scope of the claims herein.

#### 3. Traffic Rules

A rule generally includes a traffic class and a policy associated with the class. A class can have several policies that apply at different time intervals. 'Rule' is also used to refer to the policy or to a specific row in the present tool user interface. The present tool user interface is described in, for example, U.S. application Ser. No. 60/067,857 commonly assigned, which is hereby incorporated by reference for all purposes.)

#### 4. Traffic Entities

The present tool refers to entities in at least two different contexts: defining traffic classes and viewing traffic profiles. For example, a network entity generally refers to an IP address, host, subnet, IP net, IP range, URL or a group of other network entities. A service entity refers to a single service or a group of services. A native entity is referred to in viewing traffic profiles. No rule setting or configuration is required to monitor these entities. When the present tool is installed, it begins to profile traffic based upon detected services, clients, or servers, all of which are called native entities.

#### 5. Guidelines for Developing Traffic Policies

The present invention provides some guidelines for developing traffic policies. For example, to develop meaningful and effective traffic policies, the present tool may need to understand and take into account one or more of the following:

The kind of business being performed by the user over the Internet. If the user is an ISP, the user may need to develop a business/pricing model that leverages the features of the present tool. If the user is managing corporate access to the Internet, the user may want to identify any business critical services being provided over the Internet

The priority of clients, servers and URLs hosted in the user's network or servers access over the Internet. This can be organized as business critical, casual and personal.

The properties of different applications being used, whether they utilize lots of bandwidth or not. The user

may also need to account for the type of files commonly download by users or from the Web site.

Measure and analyze traffic using the present tool's profiles. Additionally, monitoring of selected entities (e.g., users, services) may also be useful.

In a further embodiment, the present tool provides some general guidelines of some commonly used applications. These guidelines should be used in conjunction with business driven priorities, traffic profiling, and selective real-time monitoring to establish an effective traffic policy. Selected guidelines are defined as follows, but are not limited to these.

Delay-sensitive low bandwidth applications, such as TELNET and DNS, are controlled best by setting a high priority policy. The present tool can give the highest priority to all network control traffic, such as QoS signaling, session establishment, domain lookup and routing protocols.

Streaming multimedia applications, such as Real Audio/Video and Vxtreme, can hog allot of bandwidth but are also delay and bandwidth sensitive. If they are not critical, they are controlled best by setting a high priority and a policy to limit admission of sessions so that bandwidth use is capped but admitted sessions have a reasonable quality.

Push technologies, such as PointCast and Marimba, download large files, are not delay or bandwidth sensitive and usually not business critical. They are best controlled by a limiting bandwidth policy and a low priority.

Bulk-data non-interactive applications, such as SMTP and NNTP, should be guaranteed a small bandwidth minimum so that they are not totally squeezed out by congestion or control policies.

Bulk-download, nominally interactive applications, such as FTP or some HTTP downloads, are commonly used in a variety of situations, ranging from critical to casual. Differentiating various types of usage in this case can usually be made only on the basis of file types and/or source or destination addresses. In this case, a small minimum can be guaranteed for more important use.

In bulk-download applications (e.g., file size > 20 K Bytes), overall congestion and burstiness can be controlled by slightly limiting this traffic, even if it is just a little below the total available bandwidth (e.g., 90%). The present tool can provide smoothing controls on this traffic without impacting overall perceptible performance for these downloads. This is particularly useful at lower link speeds (128 K and below).

Mission critical applications, such as Lotus Notes, Oracle SQLNet, and LDAP, are controlled best by setting a high priority with a guaranteed bandwidth minimum.

The above provides some guidelines for commonly used applications according to the present invention. Using the above guidelines, the present tool can effectively allocate bandwidth on a network, for example. Again, the above guidelines are merely examples and should not limit the scope of the claims herein.

In a specific embodiment, the present tool provides a comprehensive, flexible, rule-based paradigm for implementing traffic control, as illustrated by a simplified flow diagram 800 of FIG. 8. This flow diagram 800 is merely an illustration and should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Before explain-

ing the flow diagram, it may assist the reader by reviewing some general terms used herein.

These terms include, among others, "rules" and "classes" and "policies." Rules can be created for very specific groups of flows or more general groups of flows, which are commonly all the stuff that transmits to and from a link to a gateway point. Groups of flows are also referred to as traffic classes, but are not limited to such classes. Classes also can be defined by source, destination, application, file types, URLs, and other features. Policies can be specified to control traffic flows in terms of overall bandwidth guarantees, bandwidth limits, priority of service, how individual sessions within a class are serviced or admitted, and other aspects. The present tool also has intelligent policy validation that prevents users from defining any contradictory or ambiguous rules. Policy validation is generally a higher level check used by way of the present method.

The present method occurs at start, which is step 801, for example. In general, a flow of information or data or packets of information enter a gateway point, where the present tool sits. The present method classifies (step 803) the flow of information. Groups of flows can be referred to as traffic classes, but are not limited to such classes. Classes also can be defined by source, destination, application, file types, URLs, and other features. Other examples of classes were previously noted, but are not limited to these classes. In general, step 803 classifies the flow of information received into one of a plurality of predetermined classes.

The present tool measures parameters for each of the classes in step 805, which were received, for example. These parameters are based upon the policy or rule, which may be applied in a later step. As merely an example, parameters include the class itself, file sizes, and other information, which can be used by the policy or rule to apply the policy or rule to improve the quality of service for the network. After measuring the parameters, the present method applies a time stamp (step 807) on the parameters to correlate the class of information received to a time, for example.

A step of determining whether to apply a policy occurs in the next step 809. For example, if the class and the time (and the link state in some embodiments) meet predetermined settings, the policy is applied to the class in step 811 through branch 810. Alternatively, if one of the elements including the class, the time, or the link state do not meet the predetermined settings, the policy does not apply and the process continues to measure parameters through branch 808. Alternatively, the process continues to measure parameters through branch 821 after the policy is applied to the flow of information for the class.

Depending upon the application, the policy is used to improve the quality of service of the network by performing at least one of a number of functions for the class of information from the flow. These functions include, among others, bandwidth guarantees, bandwidth limits, setting priorities, admission control. The present process can also halt or stop as shown in step 815. The steps occur, in part, by way of the modules, which were previously described, but can also occur using other techniques including a combination of hardware and software, for example. These sequence of steps are merely illustrative and should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other modifications, alternatives, and variations.

The present tool controls flow of information using a combination of flow control and queuing algorithms for both outbound and inbound traffic, as shown in the simplified flow diagram 900 of FIG. 9. The algorithms work in an

define  
template

assign  
policy

monitor  
+ stream  
enforce  
policy

limit  
bandwidth  
+ access

adaptive manner to control traffic at the source taking into account existing network congestion conditions. This results in control and measurement capabilities that scale from low to high link speeds. Thousands of concurrent sessions and hundreds of rules while maximizing link utilization and minimizing congestion and packet dropping at the bottleneck point.

The present method begins with step 901, for example. In general, a flow of information or data or IP packets of information enter a gateway point, where the present tool classifies the flow of information. Groups of flows can be referred to as traffic classes, but are not limited to such classes. Classes also can be defined by source, destination, application, file types, URLs, and other features. Other examples of classes were previously noted, but are not limited to these classes. In general, step 903 classifies the flow of information received into one of a plurality of predetermined classes including for example a first group of information and a second group of information. The first group of information is generally more easily controlled using flow control techniques. The second group of information, however, is generally more easily controlled using queuing techniques. Classification occurs using an intelligent adaptation process, which separates for example TCP based 905 information and non-TCP based 907 information.

Information to be flow controlled goes to a flow control 909 engine or module. The flow control engine or module outputs 913 the information in a controlled manner using conventional flow control techniques such as those used for TCP control over the Internet. Alternatively, information to be queued goes to a queuing 911 engine or module. The queuing engine or module separates the flow of information into selected types of information, which are placed in selected queues to be transferred to a link. These selected types include RSVP, IPSEC, ICMP, and others. Each of the selected types of information are queued 915 according to type. The flow controlled information and the queued information transfer to a link or a user in a controlled manner.

The present combination of flow control and queuing can be applied to an Internet environment according to an embodiment. In particular, the application can have, for example, variation in link speeds (e.g., 14.4 kbits/second to T3 and beyond), variation in application requirements (e.g., DNS/TELNET versus bulk-data transfer applications real-time versus non-real-time applications, TCP versus UDP etc.), and variation in user requirements (e.g., casual versus business-critical). The present embodiment uses both flow control and queuing techniques, rather than relying upon a single one of these techniques, as is taught by conventional thinking. In some embodiments, the present combination of elements uses flow control techniques for a WAN access link for incoming packets of information. The present embodiment uses queuing in for outbound traffic direction, which prioritizes packets over a possible bottleneck at the WAN link. In other applications, traffic control such as queuing in the inbound direction is not used or inbound traffic is substantially free from queuing. In these embodiments, flow control can be used for inbound traffic in combination with outbound queuing techniques. Of course, the specific combination of queuing and flow control depend upon the application.

The present flow control and queuing techniques can be performed by way, at least in part, of the FAIR module, which has been described. This module is predominately software based and can be implemented on a WindowsNT™ operating system from Microsoft Corporation,

but is not limited to this operating system. Additionally, the present method can be further separated or even combined using a variety of software and hardware elements.

In a most general application, the present method provides a combination of flow control and queuing techniques for flow control of information to and from an network of computers. Flow control behavior refers to having end-points adjust their transfer rates in response to congestion indicators or under gateway control. This applies to both inbound and outbound traffic. In the preferred embodiment, the end-points implement TCP/IP flow control.

TCP flow control uses the concept of "window size" to enable a receiver end-point to indicate how much data the source end-point can send in a burst at any given time. To do this, the receiver transmits a window size limit to the source. TCP utilizes timeouts and duplicate acknowledgment signals (ACKs) to initially determine network congestion, and then utilizes the concept of window size as a tool to prevent and respond to the congestion. To do all this accurately and efficiently, TCP uses a half-dozen subtly intertwined algorithms. Congestion control is done reactively and coarsely and typically involves long delays and retransmitted traffic on the network. ACKs are used to avoid overall network collapse.

In an embodiment, a gateway at the route bottleneck (e.g., the Internet access line) is used to invoke this window size flow control mechanism in a proactive manner by screening bandwidth use and updating the sources. Typically, control applies over relatively large time scales (typical Internet round-trip times).

In alternative embodiments of the present invention, ICMP Source Quenching can also be used to serve as a flow control mechanism. ICMP Source Quenching is an IP mechanism typically invoked by routers to indicate buffer overflow. BSD based TCP/IP stacks will effect a sharp backing off of TCP data traffic bursts on receipt of such ICMP packets.

To achieve control over inbound traffic, the present invention invokes flow control at the source to spoof window size packets. This technique is also applied as a high-level control over the outbound fair scheduler in order to control congestion preferably when traffic is not subject to admission control. Further, a fair scheduler for inbound communications is used to identify which classes may need to be controlled at the source.

The present tool also supports a rule-based measurement and real-time monitoring capability for highly granular, and time-critical monitoring. For example, problems may be diagnosed by setting up a rule for measuring the affected entities and tracking key measures in a performance monitor. As it is deployed at a key access point, the present tool is in a position to monitor a problem area in the context of overall traffic quickly and easily.

In an embodiment, the present tool provides congestion, utilization, and performance degradation reports, which make day to day troubleshooting much simpler and serve to justify or validate policy setting decisions. For example, a chronic problem affecting a service through a day period (i.e., 24 hour) can be monitored by a combination of real-time monitoring, which will be described in more detail below, and congestion reports. By monitoring and using the reports, it may be determined that the affected service is not getting its due share of bandwidth, or a limitation exists with the server or in the Internet backbone.

Conclusion

In the foregoing specification, the invention has been described with reference to specific exemplary embodiments

TCP  
based  
info  
used

Real-time  
monitoring  
+  
Congestion  
reports

thereof. Many changes or modifications are readily envisioned. For example, the present invention can be applied to manage a variety of TCP/IP network traffic types for the Internet and Intranet. Further, the techniques can also be applied to Novell SPX, Xerox XNS or any protocol with a similar 'flow-control' design that utilizes windows and acknowledgment signals (similar to ACK).

Alternative embodiments of the present invention can also be applied to a 'legacy' private WAN running IP as well as native Novell protocols if there is a need. (e.g., file server and client/server traffic). Further, embodiments of the present invention can include monitoring, billing, and reporting features, thus allowing for enhanced client billing and internal cost accounting of network usage.

These techniques are preferably implemented within a firewall platform to solve the provide the following benefits: bidirectional bandwidth management of network links carrying TCP traffic; reactive (short-time scale) and proactive (long time scale) control mechanisms; and gateway (local) and end-end (global) techniques for bandwidth control. This solution reduces their contribution to congestion in the Internet; and operation in a present day heterogeneous wide area networks, such as the Internet, without requiring any client, server or router changes.

The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. It will, however, be evident that various modifications and changes may be made thereunto without departing from the broader spirit and scope of the invention as set forth in the claims.

What is claimed is:

1. A system for monitoring quality of service within a network of computers, said system comprising:
  - a network of computers, each being coupled to each other to form a local area network;
  - a link coupled to said network of computers;
  - a traffic management tool coupled to said link, said traffic management tool configured to receive a flow of information, the flow of information including TCP based information and non-TCP based information, said traffic management tool comprising a flow control module and a queuing control module, wherein the flow control module is configured to transfer said TCP based information and the queuing control module is configured to transfer said non-TCP based information.
2. The system of claim 1 further comprising a router between said link and said network of computers.
3. A method of managing information at a gateway site for improving quality of service to a network of computers, said method comprising:
  - providing a flow of information to a bandwidth management tool disposed at a server;
  - ~~classifying said flow information into at least a first portion and a second portion, the first portion including~~

removal of header from packet  
+ resending  $\Rightarrow$  TCP  
- info in header = metadata

~~TCP-based information and the second portion including non-TCP based information;~~

applying flow control to said first portion of said information to be transferred to a first user location at said local area network; and

queuing said second portion of said information to be transferred to a second user location at said local area network.

4. A system for managing information at a gateway site for improving quality of service to a network of computers, said system comprising a computer memory comprising:

- a first code directed to classifying a flow information into at least a first portion and a second portion, said first portion including TCP based information and said second portion including non-TCP based information;

- a second code directed to applying flow control to said first portion of said information to be transferred to a first user location; and

- a third code directed to queuing said second portion of said information to be transferred to a second user location.

5. The system of claim 1 wherein said link includes an incoming flow of information selected from source, destination, application, file type, or URL.

6. The system of claim 1 wherein said traffic management tool implements traffic control using one or more policies in the presence of traffic congestion or competition.

7. The system of claim 6 wherein said one of said policies is bandwidth guarantee, said bandwidth guarantee providing a selected bandwidth for a selected traffic class.

8. The system of claim 6 wherein said one of said policies is bandwidth limit, said bandwidth limit providing a bandwidth limit for a selected traffic class.

9. The system of claim 6 wherein said one of said policies is setting priorities, said setting priorities establishing an order for allocating bandwidth for selected classes.

10. The system of claim 6 wherein said one of said policies is admission control, said admission control establishing conditions for selectively admitting a service.

11. The method of claim 3 wherein said second portion of said information is selected from RSVP or IPSEC (e.g., non-TCP on top of ITP).

12. The method of claim 3 wherein said flow of information is classified IP packets.

13. The method of claim 3 wherein said first user location is the same as said second user location.

14. The method of claim 3 wherein said step of applying flow control occurs in an intelligent adaptation engine.

15. The method of claim 3 wherein said step of queuing occurs in an intelligent adaptation engine.

16. The method of claim 3 wherein said second portion of said information is substantially free from a TCP protocol.

\* \* \* \* \*